Master Seminar:
Machine Intelligence with Deep Learning
Introduction

Joseph Bethge, Christian Bartz, Mina Rezaei, Dr. Haojin Yang
Internet Technologies and Systems
Hasso Plattner Institute, University of Potsdam

Content

- **Teaching team**
- Multimedia analysis and Deep Learning
- Topic presentation
- Important information

Personal Information

Christian Bartz, M.sc

- Research background
  - 2010~2013 Bachelor Degree (Hasso-Plattner-Institute)
  - 2013~2016 Master Degree (Hasso-Plattner-Institute)
  - 2016~ PhD Student at Hasso-Plattner-Institute
- Research interests
  - Computer vision, deep learning, text recognition

Personal Information

Joseph Bethge, M.sc

- Research background
  - 2010~2013 Bachelor Degree (Hasso-Plattner-Institute)
  - 2014~2017 Master Degree (Hasso-Plattner-Institute)
  - 2017~ PhD Student at Hasso-Plattner-Institute
- Research interests
  - Computer vision, deep learning, binary neural networks

Personal Information

Mina Rezaei, M.sc

- Research background
  - 2005.10-2008.03 Azad University, Arak, Iran
    B.S c. Computer Engineering
  - 2010.10-2013.03 Shiraz University, Shiraz, Iran
    M.Sc. Artificial Intelligence
  - 2015.11-now PhD student at HPI
- Research interests
- Deep Learning for Medical Image Analysis

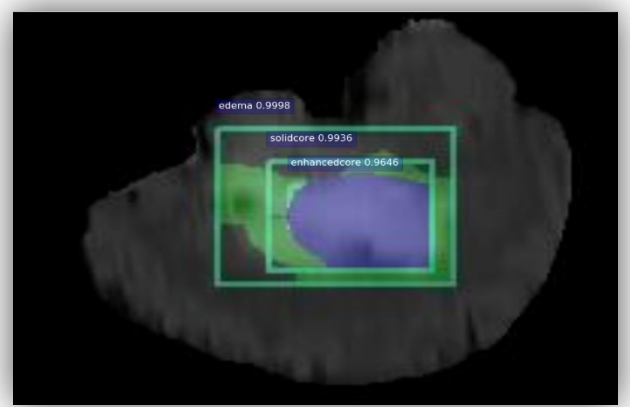

Dr. Haojin Yang

- Dipl.-Ing study at TU-Ilmenau (2002-2007)
- Software engineer (2008-2010)
- PhD student, internet technology and system HPI (2010-2013)
- Senior researcher, Multimedia and Deep Learning research team
- Research interest: multimedia analysis, computer vision, machine learning/deep learning

Research Group:

Content

- Teaching team
- **Multimedia analysis and Deep Learning**
- Topic presentation
- Important information

Content

- Teaching team
- Multimedia analysis and Deep Learning
- **Topic presentation**
- Important information

Text Localization with Deep Reinforcement Learning

- Nowadays text localization typically based on fully supervised object detectors

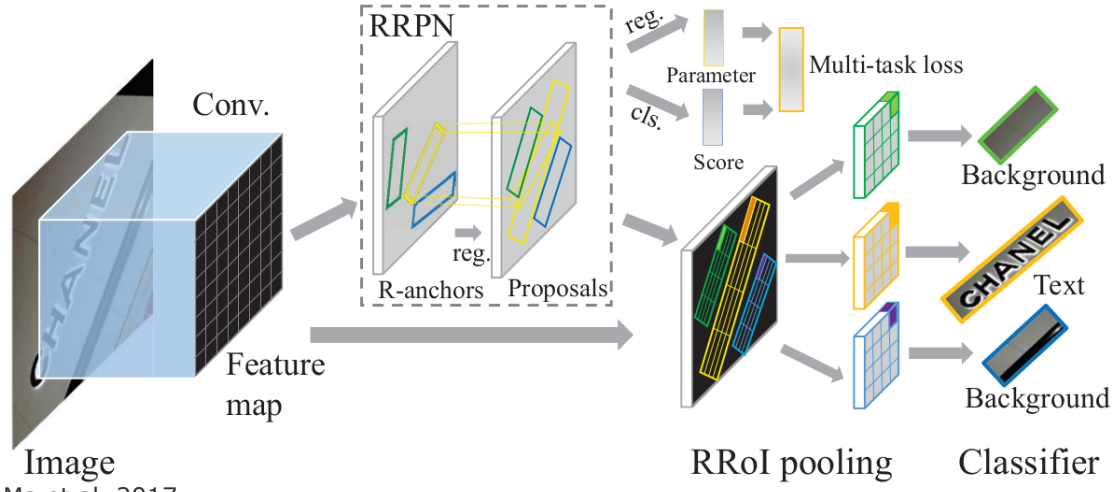

Ma et al. 2017

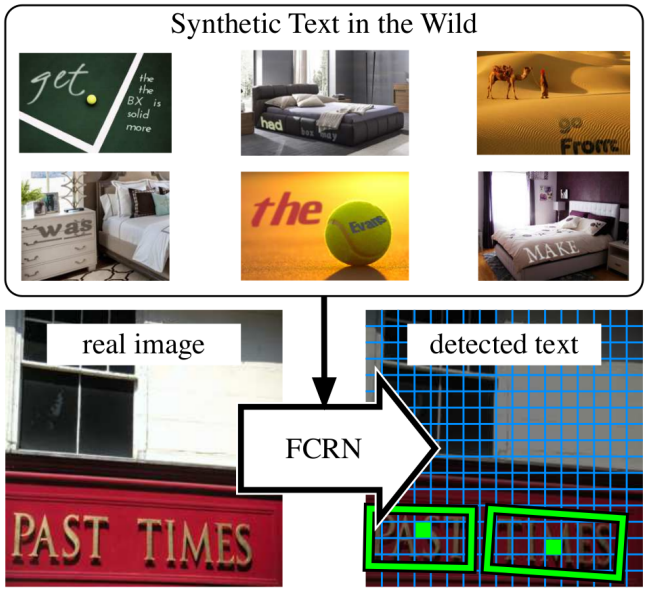

Gupta et al. 2017

Text Localization with Deep Reinforcement Learning

- How about a system that behaves like a human?

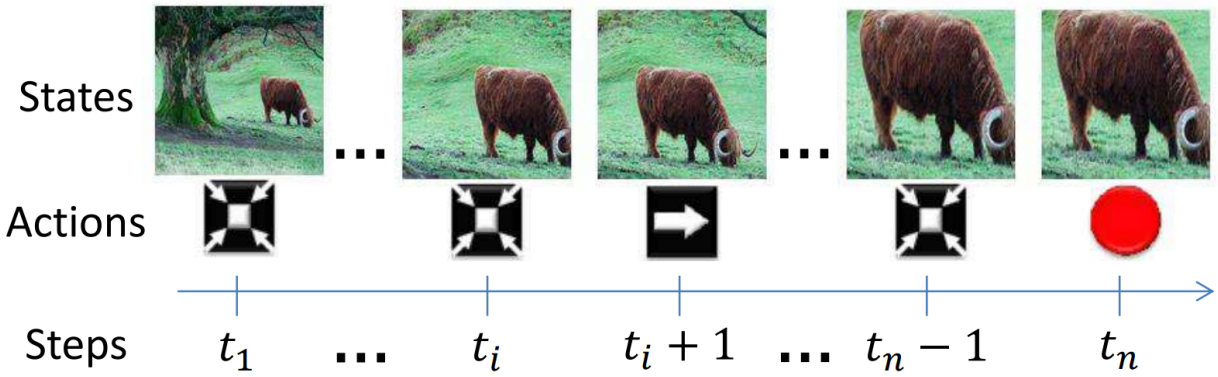

Legend:

Horizontal moves

Right Left

Aspect ratio changes

Fatter Taller

Vertical moves

Up Down

Trigger

Scale changes

Bigger Smaller

Text Localization with Deep Reinforcement Learning

- Plan:
  1. Learn about reinforcement learning
  2. Train agent for text localization
  3. …
  4. Profit!

Binary Neural Networks

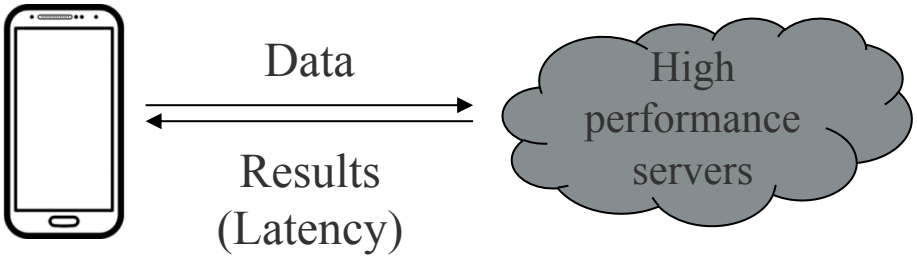

processing in the cloud

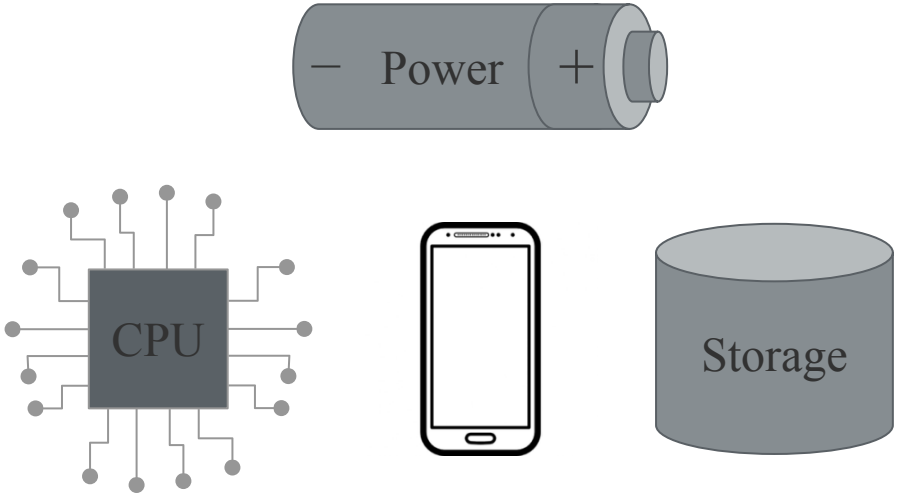

processing on device

Binary Neural Networks

- Use BMXNet “2.0” based on MXNet Gluon API (Python)
  - Dynamic computational graph, easier debugging
  - Develop an application which requires: guaranteed low latency, data privacy and/or network independency
  - Specific application is open for discussion, we have a few ideas prepared
- Deploy on a mobile device, e.g. smartphone or Raspberry Pi
  - Convert model from full-precision to binary (probably Python)
  - Update code for optimized computation to BMXNet “2.0” (C++)

Hasso Plattner Institut
IT Systems Engineering | Universität Potsdam

Interpretable Deep Models

Mina Rezaei

15.10.2018

Motivation

- DL has achieved the best performance in many domains

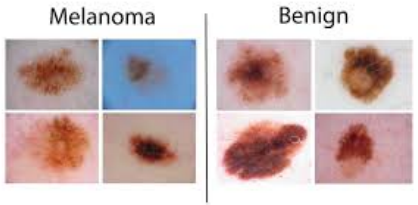

...

Black Box

100%
„No Cancer"

Source: http://interpretable-ml.org/miccai2018tutorial/

Why interpretability ?

- Verify that classifier works as we expected?
  - Wrong decisions can be costly and dangerous
- Understand weaknesses and improve classifier
- Learn new things from learning machine
- Interpretability in the sciences

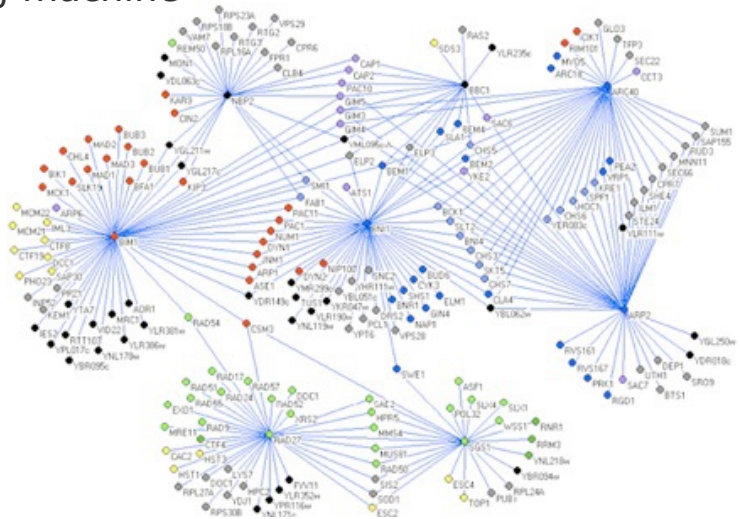

Source: http://interpretable-ml.org/miccai2018tutorial/

Dimension of Interpretability

Different dimensions of "interpretability"

prediction

"Explain why a certain pattern x has been classified in a certain way f(x)."

model

"What would a pattern belonging to a certain category typically look like according to the model."

data

"Which dimensions of the data are most relevant for the task."

Source: http://interpretable-ml.org/miccai2018tutorial/

Techniques of Interpretability

mechanistic understanding

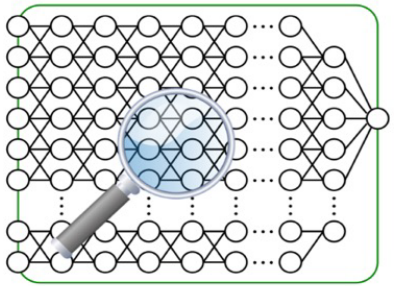

Understanding what mechanism the network uses to solve a problem or implement a function.

functional understanding

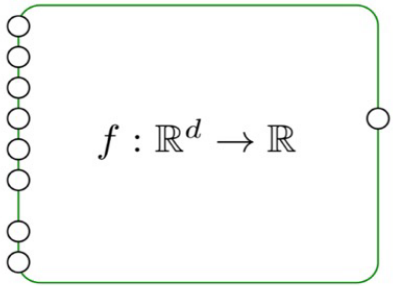

Understanding how the network relates the input to the output variables.

Source: http://interpretable-ml.org/miccai2018tutorial/

Techniques of Interpretability

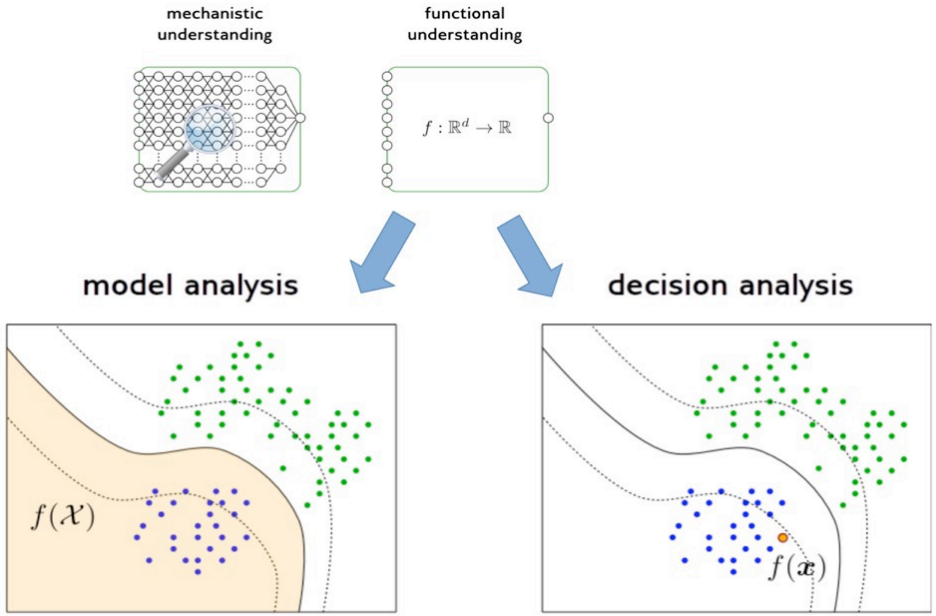

Source: http://interpretable-ml.org/miccai2018tutorial/

Model Analysis

"How does a goose typically look like according to the neural network?"

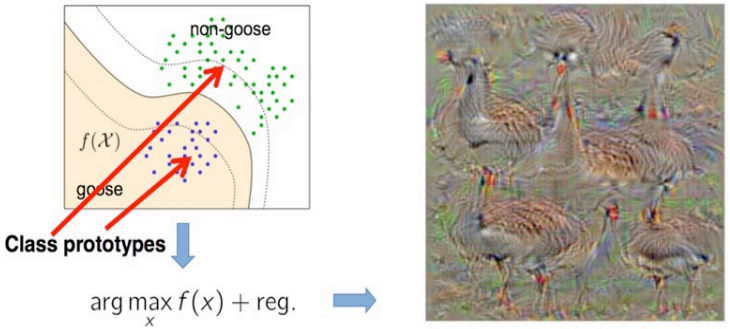

Activation Maximization

- *find prototypical example of a category*
- *find pattern maximizing activity of a neuron*

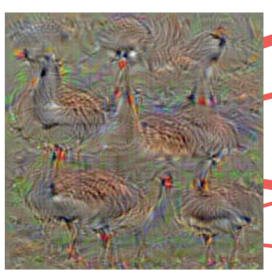

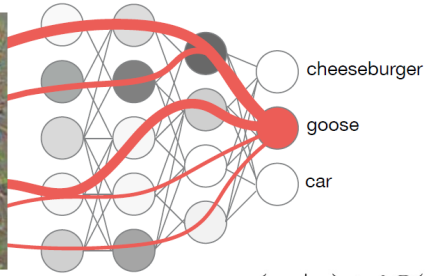

simple regularizer (Simonyan et al. 2013)

$$\max_{x \in \mathcal{X}} p_\theta(\omega_c \mid x) + \lambda \Omega(x)$$

Source: http://interpretable-ml.org/miccai2018tutorial/

Decision Analysis

- Sensitivity Analysis
- Layer-wise Relevance Propagation (LRP)
  - Heatmap of prediction

Heatmap of prediction "9"

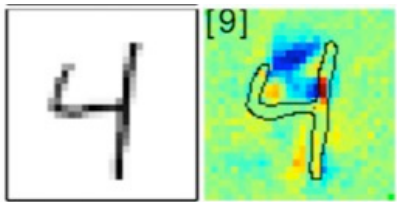

Heatmap of prediction "3"

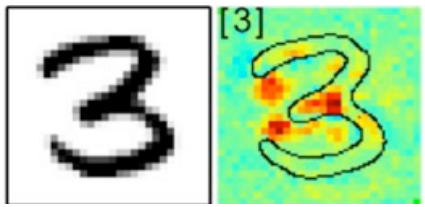

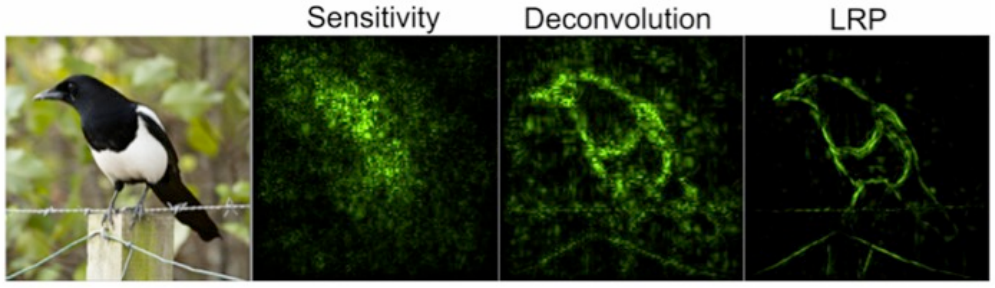

Can we objectively measure which heatmap is best ?

Source: http://interpretable-ml.org/miccai2018tutorial/

Model Analysis

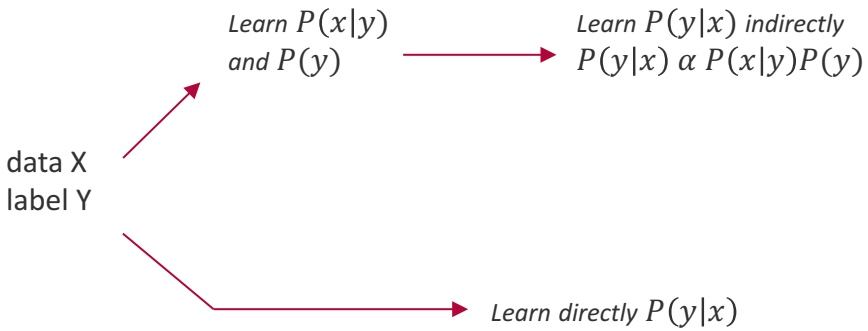

Generative Model *(G)*

Discriminative Model *(D)*

Model Analysis for Segmentation Task

Generative Model

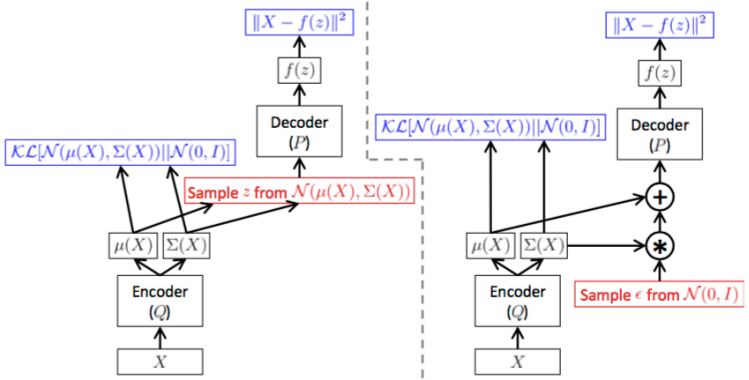

Discriminative Model

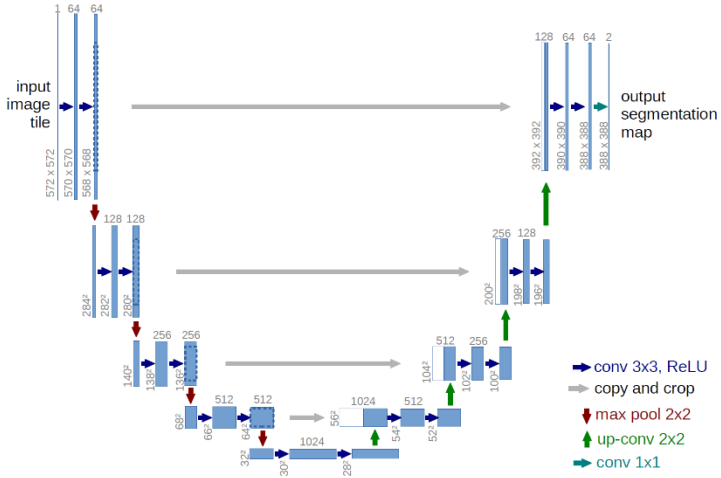

Question ?

Content

- Teaching team
- Multimedia analysis and Deep Learning
- Topic presentation
- **Important information**

Tools and Hardware

- Deep learning framework
  - Keras/Tensorflow, MXNet, Caffe/Caffe2, Chainer, PyTorch…
- GPU Servers from ITS chair

Grading Policy

- The final evaluation will be based on:
  - Initial implementation / idea presentation, 10% (03.12.2018)
  - Final presentation, 20% (04.02.2019)
  - Report/Documentation, 12-18 pages (single column), 30% (until 28.02.2018)
  - Implementation, 40% (until 28.02.2018)
  - Participation in the seminar (bonus points)

Enrollment/Anmelden

- Enroll on Doodle (link → HPI website of the course)
- Starting time: **8 a.m. 19.10.2018 (Friday)**
- Maximum number of participants: **20**

Literature

- Book: "Deep Learning", Ian Goodfellow, Yoshua Bengio and Aaron Courville, online version: www.deeplearningbook.org
- cs231n: Convolutional Neural Networks for Visual Recognition, *course of Standford University*
- Deep Learning courses at Coursera, created by Andrew Ng and deeplearning.ai, *MOOC*
- Practical Deep Learning For Coders, created by fast.ai, *MOOC*
- "Deep Learning - The Straight Dope" *http://gluon.mxnet.io, deep learning tutorials created by MXNet team*

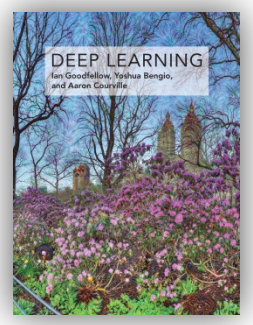

Contact

Dr. Haojin Yang

Office: H-1.22

Email: haojin.yang@hpi.de

Christian Bartz, M.sc

Office: H-1.11

Email: chrisitan.bartz@hpi.de

Mina Rezaei, M.sc

Office: H-1.22

Email: mina.rezaei@hpi.de

Joseph Bethge, M.sc

Office: H-1.21

Email: joseph.bethge@hpi.de

Thank you for your Attention!